DNSSEC Troubleshooting

Niels Monen

30-06-2011

SURFnet

Introduction

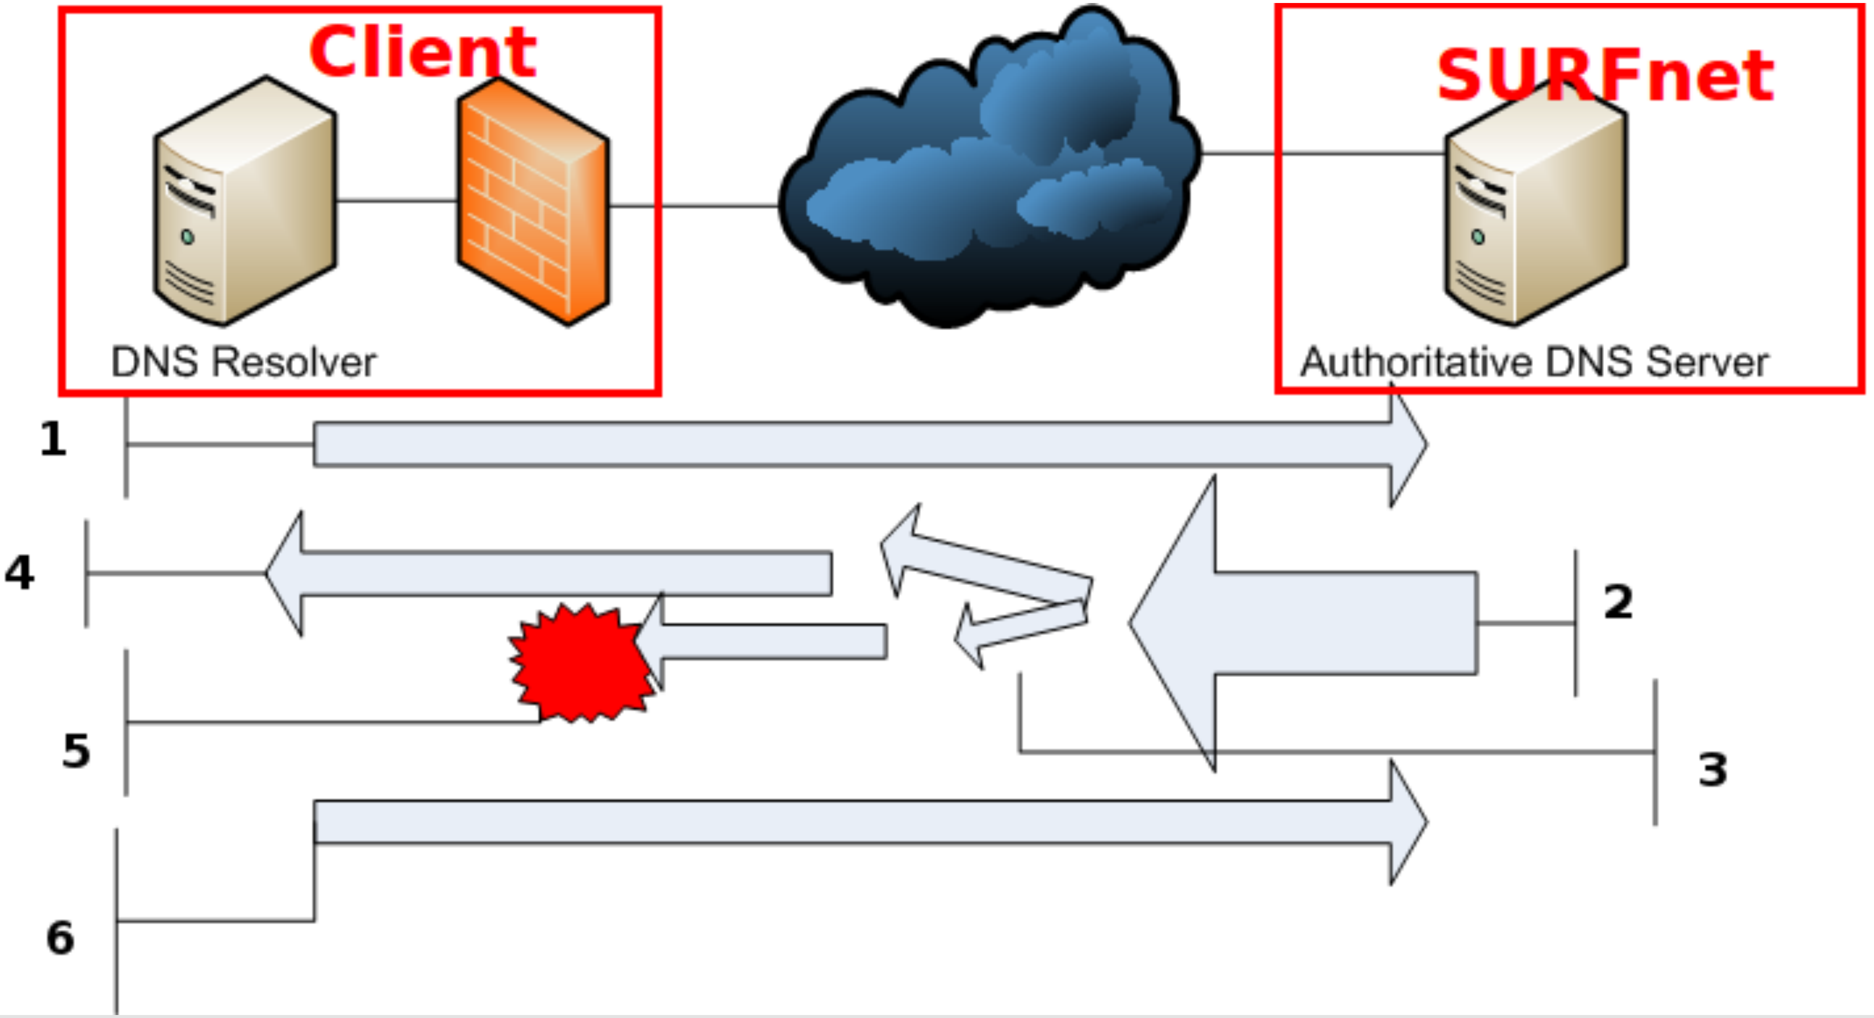

Research question

- *"Is it possible to detect if authoritative DNSSEC responses are blocked at the client side, and in particular when fragmentation occurred?"*

Sub-questions:

- *"When and where are the ICMP packets send?"*

- *"How many of SURFnet clients have this problem?"*

ICMP

- Internet Control Message Protocol
  - Typically used for error reporting in the IP layer
- RFC 792
- Many types, but for this research only Type 11
  - Time exceeded
- Code 1: Fragment Reassembly Time exceeded

DNS

- DNS (Domain Name System)
    - RFC 1035
    - UDP
    - Maximum DNS message size: 512 bytes
    - Can be extended with
        - TCP
        - EDNS0
    - DNS answer for www.surfnet.nl is only 288 bytes

DNSSEC

- DNSSEC
    - First defined in 1997 - RFC 2065
    - Latest RFC's are 4033-4035 from 2005
    - Big resource records needed
        - DNSKEY
        - RRSIG
        - NSEC(3)
        - DS
    - DNSSEC answer for www.surfnet.nl is 1659 bytes

Fragmentation

- Router receives packet PDU > next hop MTU
  - Create new IP datagrams
  - Copies Internet header into new datagrams
  - In all but last fragment, set the "More Fragments" flag
  - In all fragments, set the "Fragment Offset"
  - In last fragment, set the "More Fragments" flag to 0

Why block fragments?

- Old outdated attacks
    - Tiny fragment attack
    - Overlapping fragment attack
    - Ping of Death

Lab Setup

- Unbound
- FreeBSD IPFW
- Ubuntu 11.04 Server

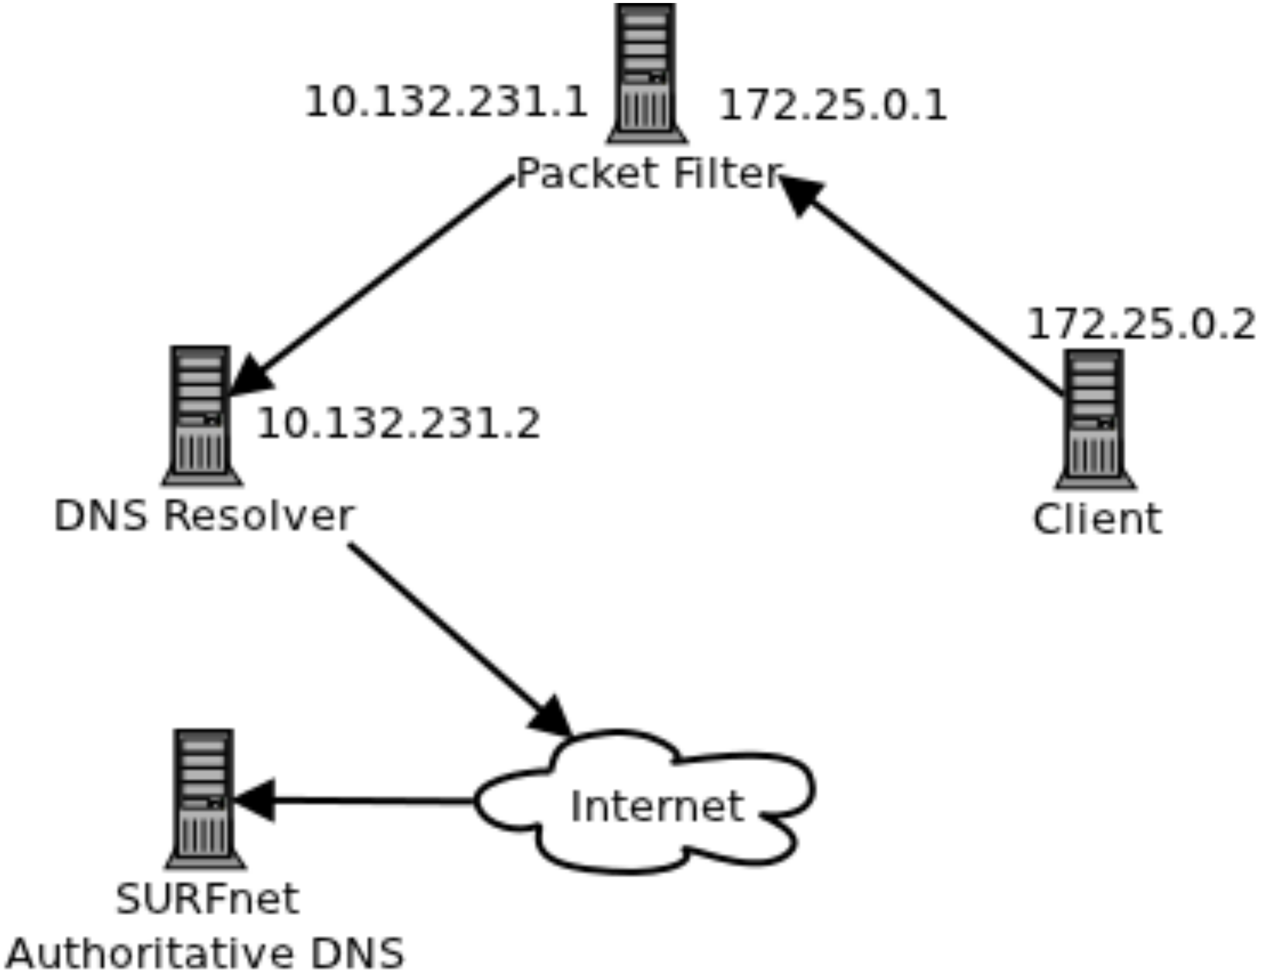

Tests (1)

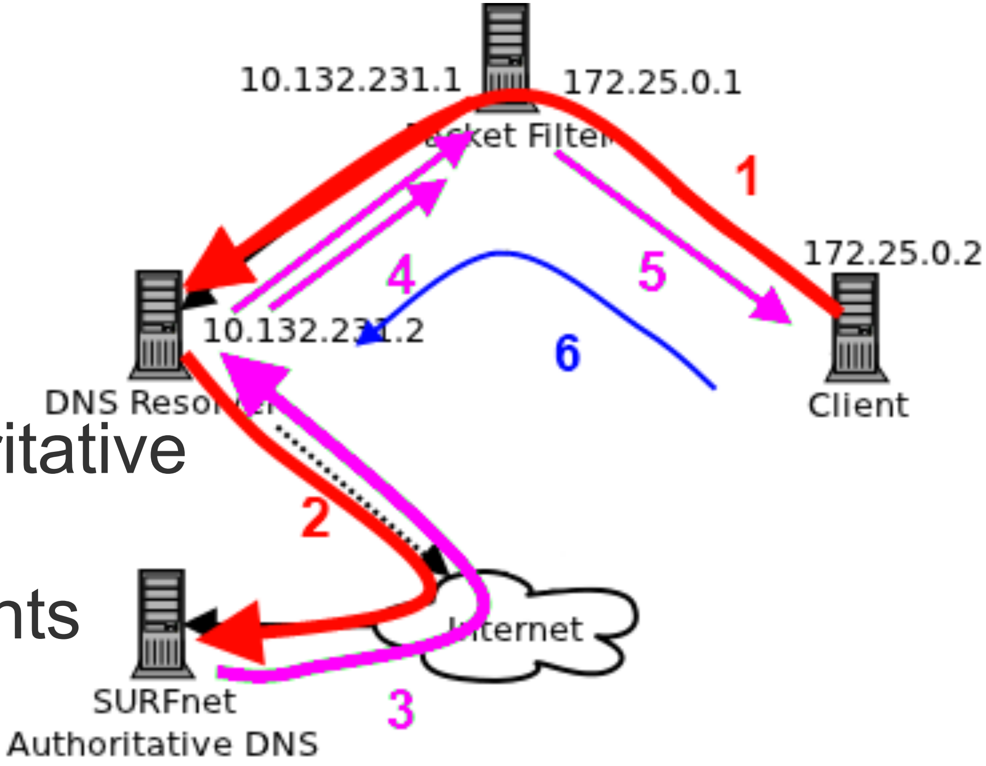

1. Request www.surfnet.nl
2. Get DNSSEC data from Authoritative
3. Send one big packet back
4. Because of MTU, send fragments
5. Only first fragment allowed
6. ICMP packet to DNS resolver

Tests (2)

- Probe
  - Written in Python + Scapy
  - Sends fragmented UDP packets to port 53
  - Checks if ICMP type 11 comes back

- Monitor
  - Extensible Ethernet Monitor (eemo)
  - Plugin to catch ICMP Type 11
  - On live environment
  - 5 Hours (12:00 - 17:00)

Results

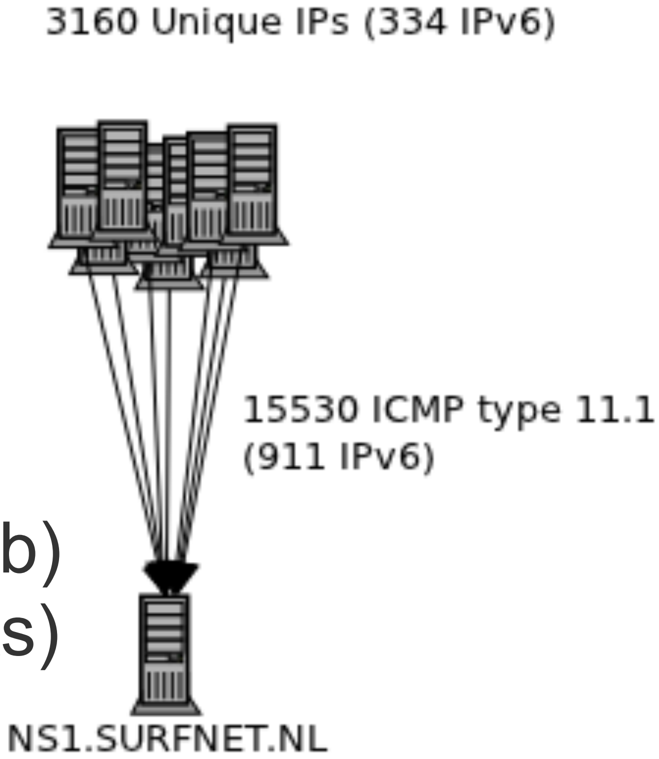

- ICMP type 11 code 1
  - Send from the client
  - Bigger than RFC 792 specified (~128b)
    - RFC 1122 (Requirements for hosts)
  - Kernel parameter when fragment reassembly times out
    - net.ipv4.ipfrag_time
  - Default is 30 seconds on modern Linux kernels
  - Default is 60 seconds on Windows 2008 R2

- 3160 SURFnet clients have this problem
  - 15530 ICMP's captured

Conclusion

"Is it possible to detect if authoritative DNSSEC responses are blocked at the client side, and in particular when fragmentation occurred?"

- It is possible by monitoring the ICMP type 11 packets
- The problem is reproducible
- At least 3160 clients have this problem
- Blocking fragments is outdated

Future Work

- Create a web page were administrators can test their servers for this problem
- Test the problem on bigger scale
- Test if ICMP packets are always arriving and correlate
- Plugin for DNS package, to monitor these ICMP packets
- Why so much IPv6?

Questions?

© Google Image Search